

UNITED STATES PATENT APPLICATION

of

Gary Robinson  
a citizen of the United States of America  
having a postal address of  
6 State Street, Suite 615  
Bangor, Maine 04401

for

**CLUSTERS FOR RAPID ARTIST-AUDIENCE MATCHING**

This United States Patent Application claims priority under 35 U.S.C. Section 119(e) from United States Provisional Patent Applications 60/165,794 filed November 16, 1999,

Gerry J. Elman, Esq. Reg. no. 24,404  
Elman & Associates  
Attorney for Applicant  
Customer No. 003775  
20 West Third Street, P.O. Box 1969  
Media, PA 19063-8969

Phone: 610-892-9580  
Fax: 610-892-9577  
email: elman@elman.com

United States of America

009775 60165794

## Clusters for Rapid Artist-Audience Matching

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent disclosure, as it appears in the PTO patent file or records, but otherwise reserves all copyrights whatsoever,

## Cross Reference to Related Applications

This application claims priority under 35 U.S.C. Section 119(e) from United States Provisional Patent Application 60/165,794, filed November 16, 1999. The entire disclosure thereof of the above-enumerated United States Provisional Patent Application, including the specification, drawings, claims, and abstract, are considered as being part of the disclosure of this application and are hereby incorporated by reference.

## Brief Summary of the Invention

The purpose of this invention is to facilitate the existence and rapid growth of a Web site (or other form of electronic service) that will distribute entertainment works to their audience more effectively than other techniques.

Online services based on this invention will

- Enable artists and entertainers to more efficiently find the consumers who will most enjoy their works
- Enable consumers to more efficiently find artists and entertainers they will enjoy
- (In some embodiments) Enable consumers of similar tastes to enjoy discussions with each other, and, in some embodiments, to interact with artists.
- (In some embodiments) Enable individuals to play an entrepreneurial role in connecting artists to their audience, wherein they may be paid for their success.

- (In some embodiments) Enable consumers and artists to enjoy the benefits of group buying: more sales to the artist and lower cost to the consumer.

## Detailed Description of the Invention

5

### Clusters: The Heart of The Invention

The centerpiece of this invention is clusters of users who have similar tastes. Users are enabled to find cluster that match their tastes, and artists are enabled to find the clusters where the users are who are likely to be interested in their work. Clusters thus serve as hubs of activity for particular tastes; in most embodiments ratings of items of interest to those tastes can be viewed, and various embodiments include various means for inter-user communication so that communities of people with similar tastes are formed.

10

Much of this disclosure will focus on music applications of the invention. However, this is merely for convenience, and applications to other fields, including works in the fields of writing and movies, fall equally within the scope of the invention.

15

### User-Created Clusters

In the some embodiment individuals are enabled to create new clusters whenever they choose. One reason for doing so is that they believe that there is a group of people which is not adequately served by any of the existing clusters – for instance, because the tastes of people of this group are substantially different from the tastes represented by any of the existing clusters.

20

The person creation a cluster will be known in this description as the “cluster initiator”.

Means are provided for cluster creators to specify a taste of the cluster (which in some embodiments is later combined with taste information from other users as described elsewhere in this document). In one embodiment, he does so by specifying ratings for various items which he

feels will be useful in defining the taste of the cluster. For example, he might give a recording of Bob Dylan's a rating of .95 (on a scale of 0 to 1) and a recording of Handel's Water Music a 0.1.

In another embodiment, he simply inputs a list of items which he feels will be the ones most-  
5 liked by members of the cluster. In many cases, these will represent his personal most-liked items. In a preferred embodiment, this list is ordered according to how well each item is liked relative to the others on the list.

Some software and network systems such as Napster and Gnutella enable file sharing, where  
10 files stored in the user's computer are made available for other users to download to their own computers. Usually, the songs a user has on his computer – and thus can make available for sharing – correspond to songs the user likes. Thus, the list of files he makes available for download can usually be presumed to represent a list of his likes, and used that way in computations. In some cases, of course, users will make songs they don't like available to other  
15 users, but some embodiments view this as happening infrequently enough that such problems may be ignored.

Then when users of the system are looking for a cluster that might be suitable for them, their tastes will be compared to that of this cluster as well as the old clusters. This new one may be the  
20 most suitable for a meaningful number of people.

In preferred embodiments, each cluster will have associated with it various facilities such as threaded discussion groups, chat, instant messaging, etc. These facilities will keep the users interested and motivate them to spend more time online. Spending this time online will provide  
25 more opportunities for advertising as well as more commitment to the cluster, increasing the probability of further online purchasing.

In some embodiments individuals are responsible for "moderating" and "administrating" these various facilities. This group may contain the cluster initiator. In many cases one person will  
30 play all these roles. For convenience, this person or group will be referred to here as the "cluster administrator."

Means are provided in such embodiments such that administrators, visitors, and members can provide reviews of items such as CD's which are in the subject domain of the service. In some embodiments, only a subset of these classes of users are enabled to input reviews; for instance, administrators and members.

5

When an item has more than one review, the usefulness of the system is enhanced by means of presenting the reviews to users in an optimal order. In the preferred embodiment, this order is based on the similarity between the user reading reviews and the user who wrote a particular review. Reviews written by authors who are closer to the reader appear earlier in the list.

10

Alternatively, the order can be dependent on the similarity of the author and the tastes of the cluster through which the item's reviews have been accessed.

15  
20  
25

In addition, in some embodiments means are provided for users to rate reviews. The system can use these ratings to determine the quality of a reviewer in general and/or the degree to which each individual user likes a particular reviewer. These factors can be used together with the similarity data, or in some embodiments, without it, to determine an ordering of reviews. For example, in one embodiment orderings are generated separately based on taste and quality. Percentile rankings are then calculated for each review. The percentiles between the two lists are averaged, and a new ordered list is created based on these average percentiles.

In some embodiments a summary of each reviewer's perceived goodness (determined by means of ratings or passive data acquisition such as measuring the average time spent reading a reviewer's reviews or the relative number of times a reviewer's reviews are sought out) is displayed.

Particular embodiments can use any means to combine the above factors. For instance, geometric or arithmetic averages can be computed, with or without weights.

30 In some embodiments, only a limited number of reviews are presented to users. These may be the ones that would have been presented first based on the ordering techniques described above.

In some embodiments where only the “best” reviews are presented, the order in which those reviews are presented may be in date-of-creation or another order such as random.

5 In some embodiments, reviews are automatically (or optionally) posted to the Usenet or other publicly-available services with links back to the cluster service or the individual cluster through which the review was written. (In some embodiments, however, there is no particular association between reviews and individual clusters; rather the association is with the item or artist being reviewed, so that all reviews are available in all clusters.)

10 In some embodiments means are provided so the administrators are paid by users for access to their clusters. In various embodiments, these payments are one-time-only, per time period (such as monthly) or per access. Reviewers can be paid similarly for access to reviews written by that reviewer. Credit card payment mechanisms or other techniques such as micropayments can be used.

15 In addition, in some embodiments facilities are provided for purchasing items directly through the site (or through partner sites reached by hyperlink). In some embodiments means are provided for users to rate their satisfaction with such purchases, and in preferred embodiments display means are provided for users to see selected summaries of these ratings, such as the  
20 average rating for the cluster (including average ratings for reviews, product purchases, general satisfaction, and any other metric of cluster success).

In embodiments where users are enabled to purchase items, preferred embodiments include means for the one or more of the cluster administrators to be paid a portion of the purchase price  
25 of the items. In various embodiments, this is a fixed percentage, a percentage that varies with volume or total revenues, or other variations.

In some embodiments advertising is used as an additional income stream.

30 In some embodiments, means are provided to enable a “group purchase” to be made wherein a number of users will purchase an item at once, thereby getting a lower price from the provider for the item. For instance, in some embodiments, means are provided for users to indicate that

they would be willing to buy an item at a particular price. When enough users are willing to buy at a particular discounted price that the provider is willing to sell the item at that price, the transaction is carried through. In other embodiments, data is stored regarding the number of people who, having purchased an item by a particular artist in the past, and/or having reviewed or rated an artist at a particular level, were willing to buy a discounted item when such a deal was presented to them. This enables the system to predict an estimate of how many people are likely to buy a particular item by the same artist if offered at a discount. This enables the administrator to purchase a substantial number of copies at the item at once, at a discount, and to pass all or part of the savings on to purchasers. In some embodiments the software is able to automatically email all members of a cluster of such deals, or to screen its output to those who have bought items from that artist or very similar artists previously.

In some embodiments, users are able to provide ratings of clusters. However, in preferred embodiments, more weight is given to ratings from people whose ratings have higher calculated "representativeness." (The concept of representativeness is discussed elsewhere in this document.)

### **Automatically-Created Clusters**

In preferred embodiments, automatically-created clusters exist instead of or in addition to user-created clusters. Note that some embodiments have the automatically-created cluster features described in this section along with a limited number of the other features of the invention which are described in this disclosure, or none of them. Automatically-created clusters have their own value independent of the other concepts described herein.

A technique for optimizing clusters based upon the principles of Shannon entropy will be described. Other techniques may be similarly applicable and also fall within the scope of the invention.

Appendix B contains instructions for creating clusters that maximize information transfer as that concept is described in the literature of Shannon entropy. The related Hartley information

approach also contains information transfer calculations, and falls within the scope of the invention, but the Shannon approach is preferred.

For completeness, Appendix C gives the Python source code to a methodology that does not use information transfer. It is based upon the standard k-means clustering concept. This example is included to illustrate the wide range of clustering approaches that fall within the scope of the invention; however, the preferred embodiment uses Shannon entropy's information transfer calculations.

This disclosure uses the term "Automatically-Created" to refer not only to systems in which clusters are created by software without manual human intervention, but also to systems in which clusters are optimized by software.

In embodiments where clusters are actually created by the software, a preferred methodology is for the administrator to set the average number of songs desired per cluster. As new songs are added to the system, new clusters are automatically created such that the average number of songs remains approximately the same; the optimization process then populates the cluster. These clusters, in various embodiments, may start out empty before they are optimized, or may be initially populated with new songs or randomly chosen songs.

In order for the software to have data to base its optimizations on, user taste data must be collected. Some embodiments do this by means of allowing users to rate songs. Preferred embodiments do this by means of passive data collection. For instances, \*.mp3 searches on the Gnutella network cause server to respond with a list of songs the user has made available for file sharing, which can be assumed, without too much error, to be a list of songs liked by that person. Radio UserLand does even better, broadcasting every song played by every user, allowing us to build a more detailed taste profile in a completely passive way. Various embodiments use various such means for data collection.

Some embodiments only allow recommendations or cluster information to be sent to processes that send realistic-seeming user data to the server. (For instance, most such embodiments would consider a process that continuously reports playing the same song to be unrealistic.)



One challenge is to associate data sent by particular processes with user identifiers that users can use to log on to Web sites. Preferred embodiments accomplish that by noting the IP address the user is accessing the Web site from, and seeing what passive data source, such as a Gnutella server, exists at the same IP address. In most such embodiments the user is then asked, via the Web interface, to confirm that he is using the type of data-broadcasting process that he is apparently using and asked whether the system has permission to link that data to his Web logon ID (or cookie, or other persistent identifier). In some embodiments, such as those involving passive data collection through Radio UserLand, a publicly available user ID for the data broadcaster is available, and that same user ID can be subsequently used by the user to log on to the Web site; the server can then easily link the data.

### ***Distributed Processing for Automatically-Created Clusters***

Preferred embodiments provide means for the computational load of the cluster calculations to be spread across more than one central processing unit.

In some embodiments, this is accomplished by having completely independent processes running on the various machines which all interact with the data as stored in a database system such as the open-source InterBase product. Each process randomly chooses a song, then finds the optimal cluster to move it to. If, when it is ready to perform the move, a check in the database indicates that another process has already moved it, then it cancels the move; otherwise it updates the database. In embodiments where more than one write needs to be performed against the database to facilitate the move, these actions are normally put into a single transaction. Using this procedure, a large number of computers can work together to perform the optimization more quickly. However, a portion of the work done will be wasted because the song in question was already moved by another process. This portion will be greater as the number of processes grows. Therefore it is preferable to have a more centrally controlled model.

Embodiments with centrally controlled models need an interprocess communication (IPC) method. Preferred embodiments use open standards such as XML-RPC and SOAP, since these enable clients to be written independently using a any of variety of programming languages. In

some such embodiments, a server process waits for registration messages from remote client processes. When a client initializes itself, it communicates with the IP address and port of the server process. The client logs on with a persistent logon ID, or the server gives it a session ID when it first makes contact. Then a portion of the workload is assigned to the client.

5

Various embodiments use various methodologies for portioning out parts of the work to the various clients. In one such embodiment, the client is sent all data needed to describe all the clusters via IPC. Then, it is assigned responsibility for a certain number of songs. It finds the best clusters for those songs. It sends that data back to the server, which subsequently updates the database. At various intervals, the cluster description data is sent again to the client, containing the results of the simultaneous work done by the various other clients.

10

In some other embodiments, only the data for a subset of the clusters is sent to the client.

Therefore, a set of clients is responsible for any particular song. Each client determines the best destination among the clusters it has the data for. Then the degree of goodness of the best choice is returned to the server by each client; the server determines "the best of the best" and updates the database accordingly.

15

One danger that arises in distributed processing embodiments is that a malicious client will be created that interacts with the server exactly as if it were a legitimate client. To avert this problem, preferred embodiments keep track of the average improvement in cluster quality per song movement. (For instance, in embodiments based on information transfer, this is based on the improvement in information transfer that occurs due to the movement.) When a client suggests a movement, the improvement associated with that movement is calculated by the server. If a client's suggestions tend to involve significantly less improvement in quality than is the norm, the system has reason to believe the client was either not written correctly or may even be malicious (trying to move songs inappropriately for the benefit of specific individuals or organizations).

20

25

The preferred embodiment accomplishes this by first storing the improvement per movement for trusted clients. These may be, for instance, clients running on the same machine as the server,

30

under control of the system administrator. As the client sends suggested movements to the server, the server determines whether the client is to be trusted.

5 In the preferred embodiment, a client's suggestions are not added to the database until the client has had a chance to prove itself. The server waits until it receives 100 suggestions. The average improvement is calculated. This average needs to be within a desired range relative to the trusted average; for instance, an installation might desire that the client's average must be within 10% of the trusted value. If it is not, that batch of 100 suggestions is thrown away. Each batch of 100 suggestions is tested separately, in case a malicious client tries to fool the server by being "nice"  
10 for a while, followed by malicious behavior.

15 Other embodiments use other techniques for screening out malicious clients. In one such technique, the value of the last 100 suggestions is averaged for each client, and the clients are subsequently ranked from least valuable to most valuable. These rankings are updated whenever new suggestions come in from a client. The last 100 suggestions from the lowest-ranking 5% (or some other number) are always ignored. Still other embodiments calculate a Bayesian estimator of the value of the next suggestion. The suggestion is counted if and only if the Bayesian estimator is within a specified range compared to the trusted client, for instance, within 1%. Other techniques are used in still further embodiments.

20 One particularly simple approach, used in some embodiments, is for the server to simply check that the preferred embodiment increases overall information transfer.

25 In some embodiments where there are so many clients that the central server does not have the processing power to check on all the clients, some clients are assigned with the task of checking that other clients are not malicious. In most such embodiments, the central server assigns these checking tasks in such a way that the assigned checking client is unknown to the possible malicious client, so there is no way for the clients to collude to trick the server.

### ***Human Input with Automatically-Created Clusters***

Preferred embodiments do not rely only on software optimization of clusters. They allow users to suggest changes. These are only made if they result in an improvement to the clustering.

- 5 For example, in one such embodiments, a Web page is made available where all the songs in a cluster are listed with checkboxes beside them. (If there are too many songs in a cluster to fit on one page, multiple pages are used. Most search engines such as Google provide fine examples of how to manage a list output when the list takes more than one page.)
- 10 There is also an entry area whether user can enter an identifier for the destination cluster. In various embodiments, an identifying number or name may be entered, or there may be a pull down list if the number of clusters is small, or a more elaborate search mechanism is used.

The user checks some of the checkboxes, specifies the destination, and indicates he is ready to continue (for instance, there may be a Continue button).

The system then determines whether the suggested movement would improve the overall clustering. For instance, in embodiments which use information transfer to measure cluster quality, the information transfer that would result if the move were completed is calculated. If it is an improvement, the transfer goes through. Otherwise, it does not, and the user is informed that the transfer didn't go through. Preferred embodiments then let the user make adjustments to his suggestion; for instance, the song listing may be presented again with the checkboxes in the state the user left them in. He can then make alterations and click Continue again.

- 25 In preferred embodiments, the user can create a new cluster as the proposed destination. It would then be initially populated with the songs he selects, if doing so increases the quality of the clustering. Many such embodiments provide a user interface whereby the user can enter songs to populate the new cluster with, without regard to there original clusters. In most such embodiments the administrator can set an upper limit on the number of clusters that may be
- 30 created in this way.

The embodiments discussed here thus join human input with computer optimization in such a way that the human input is smoothly integrated into the process. All accepted human input furthers the aim of improving the clustering.

## 5 ***Names for Automatically-Created Clusters***

Preferred embodiments provide input means for users to name automatically-created clusters.

In one such embodiment, a page is presented in which there are 20 text input areas, each providing enough space to enter a name. When a name is entered into one of the text areas (and Submit is clicked), the name may not be removed except by an administrator for a period of one week. Next to each name is a set of radio boxes labeled “no opinion, poor, fair, good, excellent”. Users can thus rate any or all of the names. User identification is carried out by means of a logon requirement, cookies, or other means; only 1 vote per user per name is allowed.

An overall rating for each name is determined by means of averaging the ratings, ignoring “no opinion” ratings.

After a name has been displayed for one week, if it is not among the top 50%, it is deleted, and any user can enter a new name.

Only one name per user at a time is accepted in the list.

At any point in time, the highest-rated name is used as the name of the cluster, displayed wherever it is convenient to display such a name. In many embodiments a cluster number is also displayed, which is constant over the life time of the cluster, and therefore may be useful when a reliable way of identifying a cluster is needed.

## **User-Cluster-Similarity**

In most embodiments, means are provided to compute a similarity between a user and a cluster.

In the some embodiments, users provide ratings that represent their tastes. In other embodiments purchase histories are used. In other embodiments, “passive” data collection such as tracking the artists and recordings that are downloaded and/or listened to can be used. In general, any source of information which captures the user’s preferences in the target domain is acceptable; this  
5 includes taking note of the music files made available for Napster, Gnutella, or other types of file sharing

In some embodiments, the “taste of the cluster,” its “taste signature,” is defined wholly by the administrator; in others it is defined by the members or even the visitors to the cluster, or by a  
10 combination thereof. The taste signature is stored in a database on the server. In some embodiments it takes the form of a list of artists or items considered to be “liked” by the cluster; in some embodiments this list is ordered with the most-liked items appearing at the top; in some embodiments ratings are associated with items and artists, for instance, on a scale from “excellent” to “poor”.

In each of these embodiments, where data from various users are combined to form the taste signature, appropriate means are used. For instance, where ratings are used, the ratings for various items and artists are averaged; in some such embodiments, a weighted average is used with the administrator having a greater weight than other users. In embodiments where ordered  
15 lists are used, means for combining include converting the lists to percentile rankings, averaging the percentile rankings for each album, and outputting a new ordered list in order of the averaged percentiles.

When a users wants to make use of the system, he usually does so by finding clusters of taste  
25 similar to theirs and, in preferred embodiments, with other positive characteristics.

In preferred embodiments, means are provided to display a list of clusters together with descriptions of each cluster supplied by the administrator. These descriptions, in various embodiments take the form of text descriptors such as “Jazz, with a focus on old-style  
30 Dixieland”, categories such as “Jazz”, “Garage Rock”, etc. or other means of communicating the center of the cluster.

In preferred embodiments, means are provided to search for clusters which users can believe they will be interested in. In embodiments where categories are provided, users can pick a category. In some embodiments where text descriptions are provided, users can search through these descriptions using standard text-retrieval techniques in order to find clusters relevant to their tastes.

In preferred embodiments, users can specify their personal tastes, and the system automatically lists clusters where the taste signature of the cluster is near to the taste signature of the user.

10 In preferred embodiments, when lists of clusters are presented based on any of the search techniques mentioned above, or other search techniques, the attributes mentioned above such as category and similarity to the user viewing the list are displayed, as may other cluster attributes which apply to the given cluster.

15 In some embodiments, "passive" data collection methods are used in matching clusters to users. These methods involve no special input of data indicating tastes.

In some such embodiments in the field of music, customizable Internet "radio" stations are associated with some or all clusters. Such stations play a mix of recordings using TCP/IP, multicasting, and/or other protocols to send streaming audio data (with additional video in some cases) to the user's computer where it is converted into sound. The recordings which are of the most interest to a cluster will tend to be played most often; the recording of least interest to the cluster, while still being "liked" by the cluster, will be played least often. Play rates can be used to tabulate ranks for items. In some embodiments, rank data is compiled for artists instead of, or  
20 in addition to, items. In most such embodiments, the administrator determines the play lists and relative frequency of playing various artists and cuts.

This rank data is then used for searching, whether acquired through manual user action or passively. In some embodiments, users input their favorite artists (or recordings, depending on  
30 the embodiment) in order of preference. In one embodiment, rank correlation is then used to find the closest matches, by computing the rank correlation for each cluster in turn and then picking the ones with the greatest level of correlation. In preferred embodiments, further processing is

done to calculate p-values relative to the rank correlations, and the p-values closest to 0 indicate the closest match. (This is preferable because p-values seamlessly incorporate the number of artists or items in common on the lists being matched, as well as the degree of similar ordering.)

5 In other embodiments, other means are used to measure taste similarities based on this data. In some embodiments, for instance, rank data is converted into "ratings" data by dividing the rankings into groups and considering the items (or artists) in the group of highest ranks to have the highest rating, the items (or artists) in the 2<sup>nd</sup>-highest group of ranks to have the second-highest rating, etc. (There are an equal number of groups of ranks to the number of ratings; for  
10 instance, if there is a 5-point rating scale, one embodiment would assign the top 20% of items [or artists] to the highest rating, the next 20% to the next highest rating, etc.) Then rating-based techniques such as those described in US Patent 5,884,282 for measuring similarity are used.

15 In some embodiments, other types of data than rank or ratings data are used. For instance, in some embodiments, simple counts of the number of items (or artists) in common on the two lists are used; a higher number means more similarity of taste. It should not be construed that this invention depends on the use of any particular type of this "taste signature" data.

20 In embodiments where we have only "presence/absence" data available, such as a Napster file list in which a particular song is either present or absent, a variety of calculations can be used. While the invention should not be construed to be limited to any particular calculations, several will be listed for purposes of example: the Ochiai, Dice, and Jaccard indices. In calculating these indices, some embodiments consider the entire list of songs to be the combination of all songs contained in either the cluster in question or the user's liked list. The presence and absence are  
25 determined corresponding to this expanded list. Some other embodiments consider the master list to be the list of songs liked by the user; other songs are ignored. Thus in such embodiments the user only has "presence" indicated; whereas the cluster will usually have a mix of presence and absence. Other embodiments do the reverse, taking the cluster's list to be the master list. Some embodiments further calculate statistical significances with respect to such indices, by making  
30 use of the statistical distribution of the used index (Snijders 1990). In all these cases a number is calculated which corresponds to the degree to which the user's list of songs matches the clusters' list of songs.



In some embodiments, passive data collection is done on the user side, in addition to, or instead of, doing so on the cluster side. In some embodiments, for example, use is made of the fact that users often have MP3, CD, streaming audio, or other types of music players on their machines.

5 Such players can be adapted by their programmers (and, in the case of open-source players, by any competent programmer) to store playback-based taste-signature data similar to that described for customizable Internet radio stations. In some embodiments this data is stored on the user's computer; in others it is stored on a central server. As noted earlier, lists of files made available for Napster, Gnutella, or other file sharing may be used. As before, rank correlation or other  
10 means, depending upon the embodiment, are used to determine the most appropriate clusters.

In some further embodiments, recommendations generated by clusters are integrated directly into the user interfaces of the users' players. For example, in some embodiments the software residing on the server is sent the playback data for a user, finds the most appropriate cluster, and  
15 sends the player software a list of the most highly-rated recordings. These recommendations are made available to the user (in one embodiment, by means of a pull-down menu; in another, by means of a scrolling list; in other embodiments, by other means) and the user can then choose the one he wants to hear. In various embodiments additional information may be included in the display, such as the name of the artist, the length of the song, etc.; in some embodiments, it is  
20 possible to click on a feature and be transported to a World Wide Web page with information on the recording.

In some embodiments, the user's player is sent the taste signature data associated with the various clusters and makes the decision about which cluster is most appropriate. This lessens  
25 privacy concerns because no complete record of a given individual's tastes needs to exist on the server to facilitate the recommendation process.

In some embodiments, the methods described here and other methods are used to measure similarities between individual users. For instance, in some embodiments these similarities are  
30 used to determine the order in which a user views reviews written by other users; the ones by users with the most similarity to the user reading the reviews are displayed first.

## Cluster Membership

In preferred embodiments, users can become members of clusters. In some embodiments, members of clusters are given special access to certain facilities like chat rooms and discussion boards. In some embodiments they are given special pricing considerations when making purchases.

In typical embodiments, cluster members are known to the system by a logon ID and password. Members can join a cluster they are visiting by indicating that they wish to join; in some embodiments this is accomplished by checking an HTML checkbox.

## Goodness List

In preferred embodiments, a goodness list is associated with each cluster. This somewhat corresponds to the top-40 song lists from decades ago.

Because a typical system might have hundreds or even thousands of clusters, the goodness list associated with each cluster will be highly targeted to particular tastes.

In some embodiments, manually entered ratings, supplied by the users, are averaged or otherwise combined to form the goodness score, and songs are listed in order of score.

In preferred embodiments, the necessary data is collected passively. In preferred embodiments, this data includes the number of times each user plays each song. Players or file sharing processes communicate their passively collected data to the server by using such common interfaces as SOAP, XML-RPC, or others.

At the time this disclosure is being written Radio UserLand broadcasts this data for its users by means of XML and any process that wants access to it can do get it by reading an XML file at a particular IP address. Radio UserLand broadcasts the time each song is played by each user; this data can be compiled to obtain a frequency of playing for each song.

Preferred embodiments use such data as follows. For each user:

- The number of times he has played each song in the last week (or during some other chosen time period) is computed. (Over the entire population, this results in one count per user per song.) Songs he has not played during that period are ignored in all following steps.
- The user's played songs are ranked with respect to one another according to the number of plays.
- A number between 0 and 1 is assigned depending on rank, in increments of  $1/N$ , where  $N$  is the number of songs played at least once by the user. The most frequently played song has a ranking of 1, the least, a rank of  $1/N$ . We will call these "unit ranks".

Then, for each song:

- The geometric mean of the unit ranks is computed. This is done by multiplying the unit ranks, and computing the  $M$ th root of the product, where  $M$  is the number of unit ranks that were multiplied. This geometric mean is considered to be the "goodness" of the song.

The number computed for each song as described above has two main advantages over other known approaches:

- Because of the ranking process, a particular user who tries to maliciously skew the process by playing a particular song an overwhelmingly huge number of times does not end up having any greater effect than another user who played the song only a little more frequently than other songs.
- By using the geometric mean to compute the goodness, the songs with the highest goodness values are the songs that most consistently achieve high play rates among users who have heard them. This consistency is important, because our aim is to create a goodness list that is very reliable. Ideally, a top-ranked song in the goodness list of a cluster will be very likely to appeal to everyone who feels an association to that cluster. Geometric means accomplish that aim.

Some embodiments take the geometric mean methodology a further step, and treat the ranks as p-values. These p-values are with respect to the null hypothesis that the song has no particular

tendency to be ranked above average compared to other songs. Then, the product of these p-values has an approximately chi-square distribution with  $2M$  degrees of freedom. So, instead of taking the  $M$ th root, we use the chi-square distribution to calculate a resultant "combined" confidence level, represented by another p-value. This resultant p-value can then be used as the goodness. Under this goodness measurement, the songs with the highest goodness would be even more reliably liked by a user with an affinity for the cluster than using the geometric mean method.

The problem with the chi-square method is that songs with a lot of people hearing them tend to generate better confidence levels, because there is more data to generate confidence from. This prejudices that goodness measure against new songs that few people have heard, even if they play the song extremely frequently.

However, in some embodiments, it is still considered worthwhile to display the chi-square-based goodness, to be as confident as possible that the top-ranked songs will be liked by nearly anyone who hears them, even though some even better newer songs will not get the attention they deserve.

In some embodiments, more than one goodness list is displayed, usually along with text describing the advantages and disadvantages of each one. For instance, once such embodiment displays the chi-square-based list with the heading "Old Reliable – You'll Be Sure To Like The Top Listed Ones Here!" and displays the geometric-mean-based ones with the heading: "Hottest of the Hot – The Top Ones Tend To Be Newer Songs Loved By Everyone Who's Heard Them!"

Some embodiments display other measures, some of which are more akin to traditional popularity measures, such as ranking the songs according to the number of people who have heard each one or the total number of plays it has received. Some embodiments display such numbers with the data restricted to users associated with the cluster; some do so over the entire population. Any combination of measures can be displayed.

In general, any measure that conveys the degree to which a song is popular or liked can be used.

These measures are often most valuable when the input data is restricted to members of the cluster for which they are being displayed. For instance, someone who loves serious, literary folk music may dislike all disco music. If for some reason he downloads a disco song and plays it once, he probably wouldn't play it again. But that should not cause the song to have a low goodness in lists that are displayed in a cluster that appeals to disco lovers.

Note that in some embodiments, there is no time window for the data to be considered by these calculations; in others older data is given less weight according to a decreasing scale, such as using half-life calculations for the data based upon the exponential distribution. (Given a chosen half-life, such as 30 days, one can compute the decay for any point in time using the exponential distribution. For our example, 30 days would have a decay of .5; days less than 30 would have decay values between 1 and .5; days greater than 30 would have decay values between .5 and 0.) This decay is an appropriate weight for the data points. If arithmetic averaging is used, the decay for each ranking is multiplied by the unit ranking. If geometric averaging is used, the unit ranking is used as a power for the ranking. Other decreasing scales may also be used. Different lists may have different scales. For instance, an "Old Reliable" list may have a window of one year, or include all relevant data ever collected, and a "Hottest of the Hot" list for the same cluster may have a window of one week.

## Radio

In some embodiments each cluster broadcasts its songs as many services on the Web broadcast songs using such formats as streaming mp3 and Real Audio. In some embodiments the administrator of a cluster can turn this feature on or off for a given cluster.

## All-You-Can-Eat Services

At the time of writing of this disclosure, many people in the Internet industry believe that a time will come when users will be able to access any song they want at any time, and either download it or play it in a streaming manner. Napster enables anyone to download any of a very great number of songs at any time for no charge, but its legality is in question because record companies and artists are not being reimbursed. It is generally expected in the industry that paid

services will shortly come into existence that give users similar benefits as those provided by Napster today, but legally. It is usually envisioned that a flat fee will be involved, akin to a monthly cable television bill. Cable TV is all-you-can-eat in the sense that for one fixed fee, the consumer gets to watch as much TV as he wants. The expected network-based music services are expected to also be all-you-can-eat in the sense that users can have access to as much music as they want for a fixed fee.

A marketplace may evolve in which artists charge such services different amounts based on their popularity. A relatively unknown artist might charge less than a better-known artist.

The service described in this disclosure can, in such a marketplace, be of use to all-you-can-eat services because the goodness measures can be used to determine who is good, regardless of the present popularity. Thus, an all-you-can-eat service can save money by marketing relatively unknown, but good, artists to its users; the more the users choose to download or listen to lesser-known artists, the more the service saves.

## Recommendations

In some cases, users will not want to have to think about clusters. They will simply want recommendations of songs.

Elsewhere in this disclosure means of measuring user-cluster-similarity are discussed. Recommendations are then made, in some embodiments, by finding the cluster(s) most similar to the user, and recommending the best songs in those clusters, according to the goodness measure used by the particular embodiment.

For instance, in some such embodiments, means such as the Ochiai presence/absence index are used to calculate a user-cluster similarity number where a higher value means more similarity, and a goodness calculation within each cluster is also performed, such as using the geometric mean of unit ranks, where a higher value means more goodness. The two numbers are then multiplied; we will call the result the recommendation priority. Recommendations are subsequently made in descending order of the recommendation priority.

If it is desired to give more weight to one of these factors, it can be taken to a power. The power can be tuned over time. One way to do that is to try different values, assigning each value for a significant period of time, such as a month. The more appropriate the power, the higher the following number will be: The average of (the recommendation priority times some passive or active measure of how much the user likes the song). For instance, in embodiments involving Radio UserLand, for each recommended song that the user has not heard before, we multiply the number of times the user actually plays it in the first week after receiving the recommendation by its recommendation priority, and compute the average of those numbers. The higher that average is, the better the weight is. After trying a number of weights over a period of time, the best one is chosen.

Other ways of combining the two numbers for calculating the recommendation priority are used in various other embodiments, such as adding them; and in still further embodiments, other methods are used, such as only picking one cluster for recommendations and then ordering them by goodness.

## Artist Tools

Items may be submitted by artists for examination by cluster administrators, possibly leading to ratings, reviews, or other consideration. In some embodiments special forms, such as Web form input, are provided for this purpose.

In preferred embodiments, means are provided to give artists some control over their “persistent reputations” as determined by ratings and reviews. In some such embodiments artists are given means to specify the clusters that may request or display reviews and ratings of their works. In further embodiments, clusters that cannot display or request reviews for an artist cannot receive submissions from him.

In order to assist artists in directing their submissions to appropriate clusters, preferred embodiments provide special tools. Preferred embodiments use taste-based searching. In one such embodiment, a form (such as a Web input form) is provided which allows an artist to list

similar artists. The clusters with most-liked-lists with the most artists in common with the artists' list are the best ones to submit to. In a further embodiment, these similar artists are listed in order of similarity. The rankings are then matched against the cluster's rankings on their ranked most-liked-lists using rank correlation. In still another embodiment, artists rate other artists regarding their similarity, and the cluster stores ratings of artist according to their perceived goodness. The scale may be for instance, a 7-point scale from "Excellent" to "Fair" in each case; although in one case similarity to a given artist is measured and in another case "goodness" seems to be measured, in fact the "goodness" measure is really similarity to the tastes of the given cluster. So the clusters with the best matches on these ratings are the ones to submit to in that embodiment.

In general, the various techniques mentioned earlier for enabling users to find appropriate clusters may also be used for artists, including deriving lists of songs from the files made available by the artist for file sharing via Napster, Gnutella, or other means, and/or using presence/absence indices. It should not be construed that this invention is limited to any particular means for taste-based searching.

In some embodiments, artists are given means to indicate that they wish to pay a particular individual to listen to, rate and/or write a review of their work. In some further embodiments, they can read the review and decide whether it is to be displayed online. In some embodiments, means are provided such as online credit card payment or checking account withdrawal through which the individual reviewer can be paid for doing the rating/review. In order to help the artist decide which user to ask for a rating and/or review, users (who may be Administrators or other users), each have information available online which would help to indicate their suitability. First, if they are members or administrators of relevant clusters, that provides a first level of filtering indicating that their tastes are probably consistent with the interests of the artist. In some embodiments, previous reviews by the user are available in one easily-accessed list. In addition, in some embodiments, if the user has entered his own ratings or explicit or implicit list of most-liked-artists, whether ordered or unordered, the artist can use his own similar information (with regard to similarity of various works to the artist's own work or simply with regard to the artist's own personal likes) to request that the system generate a calculated indicator of appropriateness, similar to that used when users are searching for appropriate clusters. In some embodiments artists can search for appropriate users using these means without consideration of clusters.



Features are provided for helping the artist made informed choices about which users to submit their items to for review. In some embodiments, artists are given input means to rate users on their satisfaction with the ratings and reviews they have paid for. Other artists can see summaries of these ratings, for instance, in one embodiment, averages of the ratings, in order to judge who to pay. (For instance, a reviewer may write a negative review but not make it public, and make it a useful critique of the work, which the artist can use in refining his work in the future; such a review might be negative but still valuable.) In some embodiments, users can set their own fees for reviewing and/or listening.

In addition, in some embodiments, a rating reliability number is calculated for users. This allows artists and other users to know how reliable a particular user's ratings are, helping artists judge whether to submit their items for rating and review by a particular user, and helping users decide which other users' ratings and reviews to read. See Appendix A for more detail.

Preferred embodiments, information is not given to the artist that will enable him to choose reviewers who only review highly. For instance, a preferred embodiment only enables artists access to each reviewer's reliability data and cluster membership. Artists will then be motivated to pick reliable reviewers, as compared to reviewers who often disagree with the majority, but they will not have a means to predict reviewers who only rate highly. Of course, in such an embodiment, an identifier for a reviewer that would enable the artist to associate him or her with particular displayed reviews would not be made available.

In a preferred embodiment, the system keeps track of the songs a user has been paid to listen to. It notes that user's relative play frequency for the songs in the weeks immediately after being paid, and the play frequencies in the cluster(s) to which the songs belong after some time has passed, for instance, 6 months, and the songs have had a chance to become known. Then, the rank correlation is calculated between the user's play frequency and the cluster's. This correlation is then used as the basis for recommending people to artists to pay to listen to their songs. To have a high correlation, the user must a) actually listen to the songs he is paid to listen to, and b) judge them similarly, relative to each other, to the way the cluster membership as a whole ultimately judges those same songs relative to each other. This embodiment is particularly appropriate in conjunction with the feature that displays songs ranked according to their average frequency of play among those who have heard the song at all (or other similar features). It

means that one user or a small number of users can be paid to hear a song, and if they like it, it will immediately be catapulted to the top of the goodness list for a cluster, encouraging still more people to listen to it, enabling good songs to become popular very quickly.

- 5 In some embodiments, artists don't have a choice regarding who they pay. Instead, the artist pays a fee, and the system decides the best people to expose the work to and/or extract ratings from. This simplifies things on a number of levels – there can be a less complicated user interface, and the artist needs to do less work. In some embodiments, artists are presented with information to the effect that, for a certain fixed fee per listener, we will pay as many people as he desires  
10 (within limits of availability) to listen. Other embodiments enable listeners to set their fees, and the system chooses based upon the fees and calculated reliability associated with each one.

Various forms of payment can be used in various embodiments. For instance, in some  
15 embodiments, money is not transferred, but instead an artist promises to make a certain number (or all) of his future recordings available to the listener for low or no cost.

- 20 In some embodiments, an “appropriate submission rating” is associated with each artist. Users rate artists with respect to the appropriateness of submitting the given item to the given user for ratings and review, which depends upon how well the item corresponds to the tastes of the user who is being to rate or review. The idea is to create a persistent record of the appropriateness of an artist's submissions in order to discourage him from “spamming” the clusters by submitting too broadly. Users can see a summary of appropriate submission ratings for the artist in question; in some embodiments this is a simple average; in others, it is a Bayesian estimator of the  
25 expected rating; in other embodiments, other summarization methods are used. Similarly, artists can see summaries of the appropriate submission ratings generated by various users; this helps them avoid submitting to users who tend to give inappropriately low ratings.

- 30 In some embodiments, there is a new songs list. This list simply lists songs that a relatively new, so that people who are interested in listening to new material can do so. This feature may appear in embodiments which do not contain any features for reimbursing those who listen to songs. In some embodiments where appropriate submission ratings are included, the songs may be listed in order of the measured appropriateness of the artist's past submissions. In further embodiments,

artists with the worst appropriateness measures may not be allowed to submit at all. Also, in some embodiments, artists who have submitted a certain number of songs in the past must have achieved a certain measured popularity if they are to be able to continue submitting. For instance, the average number of playings per week of the user's past submissions can be used; if it is below a certain point, no further submissions need be allowed. These calculations can be conducted globally or within the cluster membership. In order to keep this list from becoming too crowded, various means are used, such as always including songs for a limited, fixed set of time.

## Importance of the Administrator

In some embodiments, the administrator plays a role much like that of a radio "DJ." The administrator, sometimes called a "guide" in such embodiments, plays a role in which his own personality and tastes given high visibility. For instance, in some such embodiments, the administrator of a cluster is the only person who is enabled to provide ratings and reviews which are visible to visitors and members of the cluster. In such embodiments, administrators of different clusters compete with each other for the reputation of being the best and most reliable raters and reviewers; reliability is measured as discussed elsewhere. In further embodiments, non-administrators can provide ratings and reviews, but these are given subordinate visibility to those generated by the administrator.

## System Environment

In various embodiments, the system runs on the World-Wide-Web, client-server systems based on the TCP/IP or other communications protocols, as a multi-user program accessed by users through terminal emulators, or other technical means. In all embodiments, one or more CPU's run the system, and users are enabled to access it from remote sites through an appropriate means of communication.

## Glossary:

- Item: An article of the subject matter covered by a particular system. In various embodiments, an item can be a song, an album, a recording artist, a book, an author, a video, a director, an actor or actress, a painting, etc.
- User: A person accessing the system.
- Artist: Creator of items. For instance, the artist Herman Melville created the item “Moby Dick.”
- Cluster: A cluster is primarily defined by its taste. In various embodiments, clusters have associated facilities such as chat rooms, discussion groups, item purchase facilities, etc.
- Cluster Visitor: A user who is using the facilities of a cluster but who has not been registered with the cluster as a member.
- Cluster Member: A member has registered by indicating that he wants to join the cluster. In some embodiments, his taste is used in refining the taste of the cluster. In various embodiments members have special rights, such as the right to post to a cluster discussion group or the right to take special discounts when making purchases.
- Cluster Administrator: The person or group of people who (in some embodiments) defines the taste of the cluster, moderates chat and discussion rooms, sends notices of events, etc. In some further embodiments, the taste defined by the administrator is further refined by members and/or visitors.
- Taste of the cluster: In some embodiments, defined by the cluster administrator. In other embodiments, it is specified only by members by such means as averaging ratings for various items in the subject domain; in still other environments tastes specified by the administrator and members are combined to form the taste of the cluster. Tastes are specified and calculated as described in the text of this disclosure.

## Appendix A—Some mathematical aspects

This appendix discusses aspects of the invention that relate to certain mathematical calculations